arXiv:1611.06947v1 [cs.SI] 21 Nov 2016

Social Media as a Sensor for Censorship Detection in News Media

Rongrong Tao[1], Baojian Zhou[2], Adil Alim[2], Feng Chen[2], David Mares[3], Patrick Butler[1], Naren Ramakrishnan[1]

[1] Discovery Analytics Center, Department of Computer Science, Virginia Tech, Arlington, VA, USA
[2] Department of Computer Science, University at Albany, SUNY, Albany, NY, USA
[3] University of California at San Diego, San Diego, CA, USA
rrtao@vt.edu, {bzhou6, aalimu, fchen5}@albany.edu, dmares@ucsd.edu
pabutler@vt.edu, naren@cs.vt.edu

Abstract

Censorship in social media has been well studied and provides insight into how governments stifle freedom of expression online. Comparatively less (or no) attention has been paid to censorship in traditional media (e.g., news) using social media as a bellweather. We present a novel unsupervised approach that views social media as a sensor to detect censorship in news media wherein statistically significant differences between information published in the news media and the correlated information published in social media are automatically identified as candidate censored events. We develop a hypothesis testing framework to identify and evaluate censored clusters of keywords, and a new near-linear-time algorithm (called GraphDPD) to identify the highest scoring clusters as indicators of censorship. We outline extensive experiments on semi-synthetic data as well as real datasets (with Twitter and local news media) from Mexico and Venezuela, highlighting the capability to accurately detect real-world censorship events.

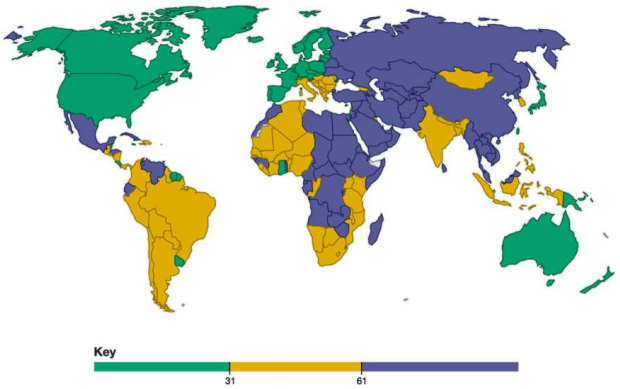

Figure 1: Worldwide freedom of the press (2014) [2]. The higher the score, the worse the press freedom status.

1. INTRODUCTION

News media censorship is generally defined as a restriction on freedom of speech to prohibit access to public information, and is taking place more than ever before. According to the Freedom of the Press Report, 40.4 percent of nations fit into the "free" category in 2003. By 2014, this global percentage fell to 32 percent [2], as shown in Figure 1. More than 200 journalists were jailed in 2014, according to the Committee to Protect Journalists. In fact, in the past three years, more than 200 journalists have been jailed annually [1].

Although the social and political aspects of news media censorship have been deeply discussed and analyzed in the field of social sciences [13, 29, 27], there is currently no efficient and effective approach to automatically detect and track such censorship events in real time.

Different from the task of Internet censorship detection in which a collection of labeled data (e.g., deleted posts or blogs in social media websites) can be collected to support supervised learning [14, 31], the detection of censorship in news media often has no labeled data available for training, and must rely on unsupervised techniques instead.

In this paper, we present a novel unsupervised approach that views social media as a sensor to detect censorship in news media wherein statistically significant differences between information published in the news media and the correlated information published in social media are automatically identified as candidate censored events.

A generalized log-likelihood ratio test (GLRT) statistic can then be formulated for hypothesis testing, and the problem of censorship detection can be cast as the maximization of the GLRT statistic over all possible clusters of keywords. We propose a near-linear-time algorithm called GraphDPD to identify the highest scoring clusters as indicators of censorship events in the local news media, and further apply randomization testing to estimate the statistical significances of these clusters.

We consider the detection of censorship in the news media of two countries, Mexico and Venezuela, and utilize Twitter as the uncensored source.

Starting in January 2012, a "Country-Withheld Content" policy has been launched by Twitter, with which governments are able to request withholding and deletion of user accounts and tweets [12]. At the same time, Twitter started to release a transparency report, which provided worldwide information and removal requests for user accounts and tweets [7]. The Transparency Report lists information and removal requests from Year 2012 to 2015 on a half-year basis. Table 1 summarizes the information and removal requests for Year 2014 on nine countries of interest. As shown in Table 1, we can see although all of these countries have ever issued account information requests, most of them did not intend

Table 1: Summary of Twitter Transparency Report for Year 2014 on nine countries of interest

Country	Account Information Request	Removal Requests	Tweets Withheld
Australia	12	0	0
Brazil	127	35	101
Colombia	8	0	0
Greece	19	0	0
Japan	480	6	43
Mexico	12	0	0
Saudi Arabia	220	0	0
Turkey	380	393	2003
Venezuela	4	0	0

to remove or withhold contents on Twitter, including Mexico and Venezuela. Therefore, we believe that Twitter can be considered as a reliable and uncensored source to detect censorship events in these two countries.

The main contributions of this paper are summarized as follows:

- **Analysis of censorship patterns between news media and Twitter.** We carried out an extensive analysis of information in Twitter deemed relevant to censored information in news media. In doing so, we make important observations that highlight the importance of our work.
- **Formulation of an unsupervised censorship detection framework**: We propose a novel hypothesis-testing-based statistical framework for detecting clusters of co-occurred keywords that demonstrate statistically significant differences between the information published in news media and the correlated information published in a uncensored source (e.g., Twitter). To the best of our knowledge, this is the first unsupervised framework for automatic detection of censorship events in news media.
- **Optimization algorithms**: The inference of our proposed framework involves the maximization of a GLRT statistic function over all clusters of co-occurrend keywords, which is hard to solve in general. We propose a novel approximation algorithm to solve this problem in nearly linear time.
- **Extensive experiments to validate the proposed techniques**: We conduct comprehensive experiments on real-world Twitter and local news articles datasets to evaluate our proposed approach. The results demonstrate that our proposed approach outperforms existing techniques in the accuracy of censorship detection. In addition, we perform case studies on the censorship patterns detected by our proposed approach and analyze the reasons behind censorship from real-world data of Mexico and Venezuela during Year 2014.

2. RELATED WORK

Here is a brief survey of three broad classes of work pertinent to our work.

Relationship between Twitter and traditional news media has been well established in many studies. Java et al. [21] studies how Twitter users report latest news on Twitter. The overlapping between Twitter and news reporting in newswire is studied in [24] and the possibility of replacing newswire with Twitter for breaking news is also explored. The role of social media in news reporting is analyzed in [30].

Event detection in social media has been studied in many recent works. Watanabe et al. [33] develops a system, which identifies tweets posted closely in time and location and determine whether they are mentions of the same event by co-occurring keywords. Ritter et al. [26] presents the first open-domain system for event extraction and an approach to classify extracted events based on latent variable models. Rozenshtein et al. [28] formulates event detection in activity network as a graph mining problem and effective greedy approaches are proposed to solve this problem. In addition to textual information, Gao et al. [19] proposes an event detection method which utilizes visual content and intrinsic correlation in social media.

Censorship is a critical problem in many countries across the world. [3] describes the situation of censorship in South Korea. Turkey, which is identified as the country issuing the largest number of censorship requests by Twitter, is studied for censorship topics by applying topic extraction and clustering on a collection of censored tweets in [31]. [17] analyzes the relationship between the Turkish government and media companies and reveals that the government exerts control over mainstream media and the flow of information. Florio et al. [18] introduces an Android app called DNSet for Turkish citizens to circumvent Internet censorship in Turkey.

3. DATA ANALYSIS

Table 2 summarizes the notation used in this work. The EMBERS project [25] provided a collection of Latin American news articles and Twitter posts. The news dataset was collected from around 6000 news agencies during Year 2014 across the world. From 4 international media sources & newspapers, we retrieved a list of top newspapers with their domain names in the target country. LNA are filtered based on the domain names in the URL links. TP was collected by randomly sampling 10% (by volume) of the Twitter data from January 1, 2014 to December 31, 2014. Retweets in TP were removed as they were not as informative as original tweets. Mexico and Venezuela were chosen as two target countries in this work since they had no censorship in Twitter (as shown in Table 1) but featured severe censorship in news media (as shown in Fig. 1).

Table 2: Description of major notations

Variable	Meaning
$\{a^t(v)\}_{t=1}^T$	time series of daily frequency of node v in uncensored Twitter dataset
$\lambda_a(v)$	expected daily frequency of node v in the Twitter dataset.
$\{b^t(v)\}_{t=1}^T$	time series of daily frequency of node v in the censored news dataset
$\lambda_b(v)$	expected daily frequency of node v in data source b
TSDF	time series of daily frequency
TP	Twitter posts dataset
LNA	local news articles dataset

3.1 Data Preprocessing

The inputs to our proposed approach are keyword co-occurrence graphs. Each node represents a keyword associated with four attributes: (1) TSDF in TP, (2) TSDF in LNA, (3) expected daily frequency in TP, and (4) expected daily frequency in LNA. Each edge represents the co-occurrence of connecting nodes in TP, or LNA, or both. However, constructing such graphs is not trivial due to data integration. One challenge is to handle the different vocabularies used in TP and LNA, with underlying distinct distributions.

To find words that behave differently in LNA comparing to TP, we only retained keywords which are mentioned in both TP and LNA. For each keyword, linear correlation between its TSDF in TP and LNA during Year 2014 is required

to be greater than a predefined threshold (e.g. 0.15) in order to guarantee the keyword is well correlated in two data sources. TSDF in TP and LNA for each node are normalized with quantile normalization. An edge is removed if its weight is less than Γ, where Γ is the threshold used to tradeoff graph sparsity and connectivity. Empirically we found $\Gamma = 10$ is an effective threshold. A keyword co-occurrence graph for a continuous time window is defined as the maximal connected component from a union of daily keyword co-occurrence graph during the time window.

3.2 Pattern Analysis

In this section, we want to answer the following question: (1) In case of no censorship in LNA, are TSDF in LNA comparable to TSDF in TP? This is important as if TSDF in LNA are always very different from TSDF in TP, we cannot make any conclusion on anomalous behaviors of LNA during a specific time period. (2) In case of censorship in LNA, are TSDF in LNA different from TSDF in TP?

To illustrate an example anomalous behavior in LNA, Fig. 2 compares TSDF in LNA and TSDF in TP during a 3-month period on a connected set of keywords sampled from data of Mexico in November 2014. We believe the example keywords are very likely to refer to the same event as the strong connectivity of these keywords, as shown in Fig. 2e, guarantees that they are mentioned together frequently in TP and LNA. The time region during which anomalous behaviors are detected is highlighted with two green markers. Since volume of TP are much larger than volume of LNA, TSDF in Fig. 2a to Fig. 2d are normalized to [0, 500] for visualization. Fig. 2a to Fig. 2d depict that TSDF in LNA fit pretty well with TSDF in TP before the highlighted time region while TSDF in LNA are significantly lower than TSDF in TP starting from the highlighted time region for all example keywords. After the disappearance of forty-three students from Ayotzinapa in the city of Iguala protesting the government's education reforms , numerous demonstrations throughout the country demanding the students' return were organized by protesters and relevant protests turned violent for the first time in early November. It is thus quite likely that the federal authorities censored key words in an effort to dampen the growing escalation of protests associated with the student disappearances. Note that our example keywords, which well described the location, population, and reasons regarding this incident, have shown a consistent pattern of abnormal behavior during our highlighted region, which starts from around November 10, 2014. Inspired from this observation, we define a **censorship pattern** exists if for a set of connected keywords, (1) their TSDF in LNA are consistently different from TSDF in TP during a time period, while (2) their TSDF in LNA are consistently similar to TSDF in TP before the time period.

4. METHODOLOGY

This section presents a novel hypothesis testing framework for characterizing the censorship patterns as discussed in Section 3 and an efficient inference algorithm for automatic detection of such censorship patterns in nearly linear time.

4.1 Problem Formulation

Suppose we have a dataset of news reports and a dataset of tweets within a shared time period in a country of interest. Each news report or tweet is represented by a set of keywords

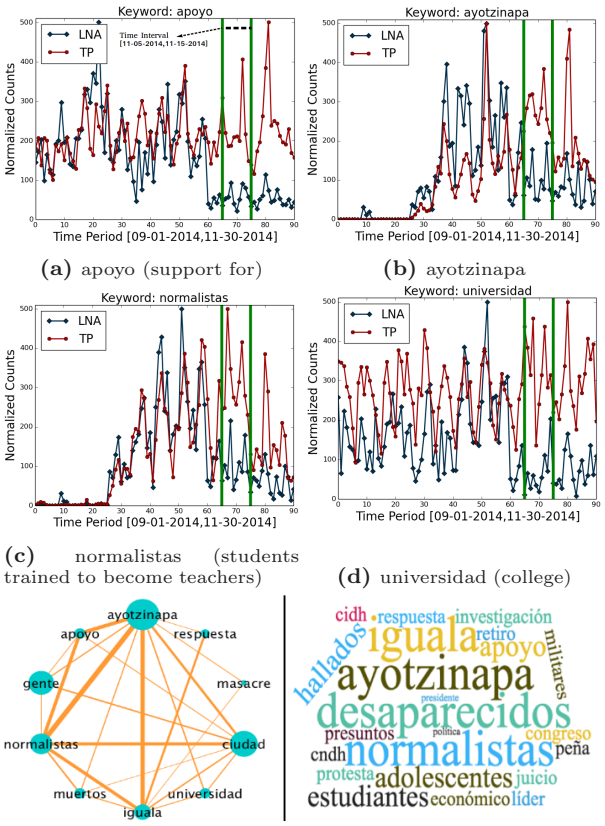

(a) apoyo (support for) **(b)** ayotzinapa **(c)** normalistas (students trained to become teachers) **(d)** universidad (college)

(e) Left: The strong connectivity of these keywords indicates their frequent co-occurrence in TP and LNA. A larger size of node indicates higher keyword frequency and a larger width of edge indicates more frequently co-occurrence; Right : word cloud representing censored keywords in LNA around 2014-11-10 in Mexico

Figure 2: Example TSDF in LNA vs. TSDF in TP for a set of connected keywords. These keywords are relevant to protests supporting for the 43 missing students from Ayotzinapa Rural Teachers' College in Mexico. We can find consistent censorship pattern in LNA starting around November 5, 2014, from when relevant protests turned violent for the first time.

and is indexed by a time stamp (e.g., day). We model the joint information of news reports and tweets using an undirected keyword co-occurrence graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$, where $\mathbb{V} = \{1, 2, \cdots, n\}$ refers to the ground set of nodes/keywords, n refers to the total number of nodes, and $\mathbb{E} \subseteq \mathbb{V} \times \mathbb{V}$ is a set of edges, in which an edge (i, j) indicates that the keywords i and j co-occur in at least one news report or tweet. Each node $v \in \mathbb{V}$ is associated with four attributes: $\{a^t(v)\}_{t=1}^T$, $\lambda_a(v)$, $\{b^t(v)\}_{t=1}^T$, and $\lambda_b(v)$ as defined in Table 2. As our study is based on the analysis of correlations between frequencies of kewyords in the news and Twitter datasets, we only consider the keywords whose frequencies in these two datasets are well correlated (with correlations above a predefined threshold 0.15). Our goal is to detect a cluster (subset) of co-occurred keywords and a time window as an indicator of censorship pattern, such that the distribution of frequencies of these keywords in the news dataset is significantly different from that in the Twitter dataset.

Suppose the chosen time granularity is day and the shared time period is $\{1, \cdots, T\}$. We consider two hypotheses: under the null (H_0), the daily frequencies of each keyword v in the news and Twitter datasets follow two different Poisson distributions with the mean parameters $\lambda_a(v)$ and $\lambda_b(v)$, re-

spectively; under the alternative ($H_1(S,R)$), there is a connected cluster S of keywords and a continuous time window $R \subseteq \{1,\cdots,T\}$, in which the daily frequencies of each keyword v in the Twitter dataset follow a Poisson with an elevated mean parameter $q_a \cdot \lambda_a(v)$, but those in the news dataset follows a Poisson with a down-scaled mean parameter $q_b \cdot \lambda_b(v)$. Formally, they can be defined as follows:

- Null hypothesis H_0:
$$a^t(v) \sim \text{Pos}(\lambda_a(v)), \forall v \in \mathbb{V}, t \in \{1,\cdots,T\}$$
$$b^t(v) \sim \text{Pos}(\lambda_b(v)), \forall v \in \mathbb{V}, t \in \{1,\cdots,T\}$$
- Alternative hypothesis $H_1(S,R)$:
$$a^t(v) \sim \text{Pos}(q_a \cdot \lambda_a(v)) \ , \ b^t(v) \sim \text{Pos}(q_b \cdot \lambda_b(v)), \forall v \in S, t \in R$$
$$a^t(v) \sim \text{Pos}(\lambda_a(v)), \qquad b^t(v) \sim \text{Pos}(\lambda_b(v)), \forall v \notin S \text{ or } t \notin R$$

where $q_a > 1, q_b < 1$, $S \subseteq \mathbb{V}$, the subgraph induced by S (denoted as $\mathbb{G}_S$) must be connected to ensure that these keywords are semantically related, and $R \subseteq \{1,2,\cdots,T\}$ is a continuous time window defined as $\{i, i+1, \cdots, j\}, 1 \leq i \leq j \leq T$. Given the Poisson probability mass function denoted as $p(x;\lambda) = \lambda^x e^{-\lambda}/x!$, a generalized log likelihood ratio test (GLRT) statistic can then be defined to compare these two hypotheses, and has the form:

$$F(S,R) = \log \frac{\max_{q_a>1} \prod_{t\in R}\prod_{v\in S} p(a^t(v); q_a\lambda_a(v))}{\prod_{t\in R}\prod_{v\in S} p(a^t(v);\lambda_a(v))} + \log \frac{\max_{q_b<1} \prod_{t\in R}\prod_{v\in S} p(b^t(v); q_b\lambda_b(v))}{\prod_{t\in R}\prod_{v\in S} p(b^t(v);\lambda_b(v))}. \quad (1)$$

In order to maximize the GLRT statistic, we need to obtain the maximum likelihood estimates of q_a and q_b, which we set $\partial F(S,R)/\partial q_a = 0$ and $\partial F(S,R)/\partial q_b = 0$, respectively and get the best estimate $\hat{q_a} = C_a/B_a$ of q_a and $\hat{q_b} = C_b/B_b$ of q_b where $C_a = \sum_{v\in S, t\in R} a^t(v)$, $C_b = \sum_{v\in S, t\in R} b^t(v)$, $B_a = \sum_{v\in S, t\in R} \lambda_a(v)$, $B_b = \sum_{v\in S, t\in R} \lambda_b(v)$. Substituting q_a and q_b with the best estimations $\hat{q_a}$ and $\hat{q_b}$, we obtain the parametric form of the GLRT statistic as follows:

$$F(S,R) = \Big(C_a \log \frac{C_a}{B_a} + B_a - C_a\Big) + \Big(C_b \log \frac{C_b}{B_b} + B_b - C_b\Big) \quad (2)$$

Given the GLRT statistic $F(S,R)$, the problem of censorship detection can be reformulated as Problem 1 that is composed of two major components: 1) **Highest scoring clusters detection.** The highest scoring clusters are identified by maximizing the GLRT statistic $F(S,R)$ over all possible clusters of keywords and time windows; 2) **Statistical significance analysis.** The empirical p-values of the identified clusters are estimated via a randomization testing procedure [23], and are returned as significant indicators of censorship patterns in the news dataset, if their p-values are below a predefined significance level (e.g., 0.05).

Problem 1. ***(GLRT Optimization Problem)*** *Given a keyword co-occurrence graph $\mathbb{G}(\mathbb{V},\mathbb{E})$ and a predefined significant level α, the* GLRT *optimization problem is to find the set of highest scoring and significant clusters $\mathbb{O}$. Each cluster in $\mathbb{O}$ is denoted as a specific pair of connected subset of keywords ($S_i \subseteq \mathbb{V}$) and continuous time window ($R_i \subseteq \{1,\cdots,T\}$), in which S_i is the highest scoring subset within the time window R_i:*

$$\max_{S\subseteq\mathbb{V}} F(S,R_i) \text{ s.t. } S \text{ is connected}, \quad (3)$$

and is significant with respect to the significance level α.

4.2 GraphDPD Algorithm

Our proposed algorithm GraphDPD decomposes Problem 1 into a set of sub-problems, each of which has a fixed continuous time window, as shown in Algorithm 1. For each specific day i (the first day of time window R in Line 6) and each specific day j (the last day of time window R of Line 6), we solve the sub-problem (Line 7) with this specific $R = \{i, i+1, \cdots, j\}$ using Relaxed-GrapMP algorithm which will be elaborated later. For each connected subset of keywords S returned by Relaxed-GraphMP, its p-value is estimated by randomization test procedure [23](Line 8). The pair (S,R) will be added into the result set $\mathbb{O}$ (Line 9) if its empirical p-value is less than a predefined significance level α (e.g., 0.05). The procedure getPValue in Line 8 refers to a randomization testing procedure based on the input graph $\mathbb{G}$ to calculate the empirical p-value of the pair (S,R) [23]. Finally, we return the set $\mathbb{O}$ of signifiant clusters as indicators of censorship events in the news data set.

Algorithm 1 GraphDPD

```
1: Input: Graph Instance 𝔾 and significant level α;
2: Output: set of anomalous connected subgraphs 𝕆;
3: 𝕆 ← ∅;
4: for i ∈ {1,···,T} do
5:     for j ∈ {i+1,···,T} do
6:         R ← {i, i+1, ···, j} ; // time window R
7:         S ← Relaxed-GraphMP(𝔾, R);
8:         if getPValue(𝔾, S, R) ≤ α then
9:             𝕆 ← 𝕆 ∪ (S, R);
10:        end if
11:    end for
12: end for
13: return 𝕆;
```

Line 7 in Algorithm 1 aims to solve an instance of Problem 1 given a specific time window R, which is a set optimization problem subject to a connectivity constraint. Tung-Wei et. al. [22] proposed an approach for maximizing submodular set function subject to a connectivity constraint on graphs. However, our objective function $F(S,R)$ is non-submodular as shown in Theorem 1 and this approach is not applicable here.

Theorem 1. *Given a specific window R, our objective function $F(S,R)$ defined in (2) is non-submodular.*

Proof. This can be proved by a counter example. ■

We propose a novel algorithm named Relaxed-GraphMP to approximately solve Problem 1 in nearly linear time with respect to the total number of nodes in the graph. We first transform the GLRT statistic in Equation(2) to a vector form. Let $\mathbf{x}$ be an n-dimensional vector $(x_1, x_2, \cdots, x_n)^\mathsf{T}$, where $x_i \in \{0,1\}$ and $x_i = 1$ if $i \in S$, $x_i = 0$ otherwise. We define $\mathcal{P}, \mathcal{Q}, \Lambda_a, \Lambda_b$ as follows:

$$\mathcal{P} = \left[\sum_{t\in R} a^t(1), \cdots, \sum_{t\in R} a^t(n)\right]^\mathsf{T}, \Lambda_a = [\lambda_a(1), \cdots, \lambda_a(n)]^\mathsf{T},$$

$$\mathcal{Q} = \left[\sum_{t\in R} b^t(1), \cdots, \sum_{t\in R} b^t(n)\right]^\mathsf{T}, \Lambda_b = [\lambda_b(1), \cdots, \lambda_b(n)]^\mathsf{T}.$$

Therefore, C_a, C_b, B_a, and B_b in Equation(2) can be reformulated as follows:

$$C_a = \mathcal{P}^\mathsf{T}\mathbf{x}, \quad C_b = \mathcal{Q}^\mathsf{T}\mathbf{x}, \quad B_a = |R|\Lambda_a{}^\mathsf{T}\mathbf{x}, \quad B_b = |R|\Lambda_b{}^\mathsf{T}\mathbf{x}$$

Hence, F can be reformulated as a relaxed function $\hat{F}$:

$$\hat{F}(\mathbf{x}, R) = \mathcal{P}^\mathsf{T}\mathbf{x} \log \frac{\mathcal{P}^\mathsf{T}\mathbf{x}}{|R|{\Lambda_a}^\mathsf{T}\mathbf{x}} + |R|\Lambda_a{}^\mathsf{T}\mathbf{x} - \mathcal{P}^\mathsf{T}\mathbf{x} + \mathcal{Q}^\mathsf{T}\mathbf{x} \log \frac{\mathcal{Q}^\mathsf{T}\mathbf{x}}{|R|{\Lambda_b}^\mathsf{T}\mathbf{x}} + |R|\Lambda_b{}^\mathsf{T}\mathbf{x} - \mathcal{Q}^\mathsf{T}\mathbf{x} \quad (4)$$

We relax the discrete domain $\{0,1\}^n$ of S to the continuous domain $[0,1]^n$ of $\mathbf{x}$, and obtain the relaxed version of Problem 1 as described in Problem 2.

Problem 2. ***Relaxed* GLRT *Optimization Problem*** *Let $\hat{F}$ be a continuous surrogate function of F that is defined on the relaxed domain $[0,1]^n$ and is identical to $F(S,R)$ on the discrete domain $\{0,1\}^n$. The relaxed form of* **GLRT *Optimization Problem*** *is defined the same as the* GLRT *optimization problem, except that, for each pair (S_i, R_i) in $\mathbb{O}$, the subset of keywords S_i is identified by solving the following problem with $S_i = supp(\hat{\mathbf{x}})$:*

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}\in[0,1]^n} \hat{F}(\mathbf{x}, R_i) \quad s.t. \quad supp(\mathbf{x}) \text{ is connected.}$$

where $supp(\mathbf{x}) = \{i | x_i \neq 0\}$ is the support of $\mathbf{x}$. The gradient of $\hat{F}(\mathbf{x}, R)$ has the form:

$$\frac{\partial \hat{F}(\mathbf{x}, R)}{\partial \mathbf{x}} = \log \frac{\mathcal{P}^\mathsf{T}\mathbf{x}}{|R|{\Lambda_a}^\mathsf{T}\mathbf{x}}\mathcal{P} + \Big(|R| - \frac{\mathcal{P}^\mathsf{T}\mathbf{x}}{{\Lambda_a}^\mathsf{T}\mathbf{x}}\Big)\Lambda_a + \log \frac{\mathcal{Q}^\mathsf{T}\mathbf{x}}{|R|{\Lambda_b}^\mathsf{T}\mathbf{x}}\mathcal{Q} + \Big(|R| - \frac{\mathcal{Q}^\mathsf{T}\mathbf{x}}{{\Lambda_b}^\mathsf{T}\mathbf{x}}\Big)\Lambda_b \quad (5)$$

Algorithm 2 RELAXED-GRAPHMP

1: **Input**: Graph instance $\mathbb{G}$, continous time window R;
2: **Output**: the co-occurrence subgraph $\mathbb{G}_S$;
3: $i \leftarrow 0$; $\mathbf{x}^i \leftarrow$ an initial vector;
4: **repeat**
5: $\quad \nabla\hat{F}(\mathbf{x}^i, R) \leftarrow \frac{\partial \hat{F}(\mathbf{x}^i, R)}{\partial \mathbf{x}^i}$ by Equation (5) ;
6: $\quad \mathbf{g} \leftarrow \mathbf{Head}(\nabla\hat{F}(\mathbf{x}^i, R), \mathbb{G})$; // *Head projection step*
7: $\quad \Omega \leftarrow \text{supp}(\mathbf{g}) \cup \text{supp}(\mathbf{x}^i)$;
8: $\quad \mathbf{b} \leftarrow \arg\max_{\mathbf{x}\in[0,1]^n} \hat{F}(\mathbf{x}, R)$ *s.t.* $\text{supp}(\mathbf{x}) \subseteq \Omega$;
9: $\quad \mathbf{x}^{i+1} \leftarrow \mathbf{Tail}(\mathbf{b}, \mathbb{G})$; // *Tail projection step*
10: $\quad i \leftarrow i+1, S \leftarrow \text{supp}(\mathbf{x}^i)$;
11: **until** halting condition holds;
12: **return** (S, R);

Our proposed algorithm RELAXED-GRAPHMP decomposes Problem 2 into two sub-problems that are easier to solve: 1) a single utility maximization problem that is independent of the connectivity constraint; and 2) head projection and tail projection problems [20] subject to connectivity constraints. We call our method RELAXED-GRAPHMP which is analogous to GRAPHMP proposed by Chen et al. [16]. The high level of RELAXED-GRAPHMP is shown in Algorithm 2. It contains 4 main steps as described below.

- **Step 1:** Compute the gradient of relaxed GLRT problem (Line 5). The calculated gradient is $\nabla\hat{F}(\mathbf{x}^i, R)$. Intuitively, it maximizes this gradient with connectivity constraint that will be solved in next step.
- **Step 2:** Compute the head projection (Line 6). This step is to find a vector $\mathbf{g}$ so that the corresponding subset $\text{supp}(\mathbf{g})$ can maximize the norm of the projection of gradient $\nabla\hat{F}(\mathbf{x}^i, R)$ (See details in [20]).
- **Step 3:** Solve the maximization problem without connectivity constraint. This step (Line 7,8) solves the maximization problem subject to the $\text{supp}(\mathbf{x}) \subseteq \Omega$, where Ω is the union of the support of the previous solution $\text{supp}(\mathbf{x}^i)$ with the result of head projection $\text{supp}(\mathbf{g})$ (Line 7). A gradient ascent based method is proposed to solve this problem. Details is not shown here due to space limit.
- **Step 4:** Compute the tail projection (Line 9). This final step is to find a subgraph $\mathbb{G}_S$ so that $\mathbf{b}_S$ is close to $\mathbf{b}$ but with connectivity constraint. This tail projection guarantees to find a subgraph $\mathbb{G}_S$ with constant approximation guarantee (See details in [20]).
- **Halting:** The algorithm terminates when the condition holds. Our algorithm returns a connected subgraph $\mathbb{G}_S$ where the connectivity of $\mathbb{G}_S$ is guaranteed by Step 4.

Time Complexity Analysis: The GRAPHDPD algorithm is efficient as its time complexity is proportional to the total number of continous time windows T^2. Therefore, the time complexity of GRAPHDPD is mainly dependent on the run time of RELAXED-GRAPHMP. We give the detailed time complexity analysis in Theroem 2.

Theorem 2. GRAPHDPD *runs in $O(T^2 \cdot t(nT + nl + |\mathbb{E}|\log^3 n))$ time, where T is the maximal time window size, nT is the time complexity of Line 5 in Algorithm 1, nl is the run time of Line 8 using gradient ascent, $|\mathbb{E}|\log^3 n$ is the total run time of head projection and tail projection algorithms, and t is the total number of iterations needed in Algorithm 2.*

Proof. As the maximal time window in input graph $\mathbb{G}$ is T, GRAPHDPD needs $O(T^2)$ iterations in its inner loop and outer loop (From Line 4 to Line 11 in Algorithm 1) to execute RELAXED-GRAPHMP (Line 7). Suppose RELAXED-GRAPHMP needs t iterations, the time complexity of each iteration has three parts: 1). $O(nT)$ the run time for calculating gradient in Line 5 of Algorithm 2; 2). $O(nl)$: the run time of Line 8 using gradient ascent where l is the number of iterations in gradient ascent method; and 3). $O(|\mathbb{E}|\log^3 n)$: the run time of head and tail projection in Line 6 and Line 9 of Algorithm 2. Hence the time complexity of RELAXED-GRAPHMP is $t(nT + nl + |\mathbb{E}|\log^3 n)$. Therefore, the total time complexity of GRAPHDPD immediately follows. As observed in our experiments, the numbers of iterations, including t and l, scale constant with respect to n, and the overall time complexity of GRAPHDPD hence scales nearly linear with respect to n. ■

5. EXPERIMENTS

Through experiments, we (1) evaluate the performance of our proposed approach in censorship pattern detection compared with baseline methods, and (2) perform case studies that analyze the censorship patterns we have found in real data. The code and datasets will be available to researchers for evaluation purposes.

5.1 Experimental Design

Real world datasets: Table 3 gives a detailed description of real-world data we used in this work. Details of collecting Twitter and news data sets are provided in Section 3. Daily keyword co-occurrence graphs, which integrate LNA with TP, are generated as described in Section 3.1.

Country	Tweets	News	# of Nodes	# of Edges
Mexico	249,124	863	3,369	93,919
Venezuela	222,072	169	2,469	37,740

Table 3: Real-world dataset used in this work. Tweets: average number of daily tweets. News: average number of daily domestic news articles. Number of nodes and edges are averaged over daily keyword co-occurrence graphs.

Data Preprocessing: The preprocessing of the real world datasets is discussed in details in Section 3.1. In particular, we considered keywords whose day by day frequencies in news media and Twitter data have linear correaltions above 0.15, in order to filter noisy keywords.

Semi-synthetic datasets: We create semi-synthetic data sets by using the coordinates from real-world datasets and injecting anomalies. Ten daily keyword co-occurrence graphs are randomly selected to inject with random true anomaly connected subgraphs using random walk algorithm [32] with a restart probability of 0.1. The number of nodes in true anomaly subgraph is x percentage of number of nodes in the daily co-occurrence graph, where $x \in \{0.05, 0.1, 0.15\}$. For convenience but without loss of generality, we specified $q_t \cdot q_n = 1.0$, where q_t controls the scale of anomaly in tweets and q_n controls scale of anomaly in local news articles. In our experiments, we set $q_t = \{1.0, 2.0, \cdots, 10.0, 15.0, \cdots, 35.0\}$, and set $q_n = 1/q_t$ correspondingly.

Our proposed Graph-DPD and baseline methods: The maximal window size T and result threshold α in GraphDPD are set as 7 and 0.05 respectively. However, our algorithm is not sensitive to the setting of T and α. We compare our proposed method with the two state-of-art baseline methods, which are designed specifically for connected anomalous subgraph detection, namely, EventTree [28] and NPHGS [15]. Model parameters are tuned following the original papers. Specifically, for EventTree we tested $\lambda = \{0.0001, 0.0006, \cdots, 0.001, 0.006, \cdots, 0.010, 0.015, \cdots, 0.1, 0.5, 1.0, \cdots, 20.0\}$. For NPHGS, we set the number of seed entities $K = 400$ and typical significance levels $\alpha_{max} = 0.15$ as the authors suggested. Since the baseline methods are designed to detect anomalies on one data source at one time, they are tested separately on TP and LNA, which are labeled as $\text{EventTree}_{\text{News}}$, $\text{EventTree}_{\text{Twitter}}$, $\text{NPHGS}_{\text{News}}$ and $\text{NPHGS}_{\text{Twitter}}$. Sepecifically, $\text{EventTree}_{\text{Twitter}}$ and $\text{NPHGS}_{\text{Twitter}}$ are burst detection baseline methods while $\text{EventTree}_{\text{News}}$ and $\text{NPHGS}_{\text{News}}$ are absenteeism detection baseline methods by some transformations on attributes.

Performance Metrics: The performance metrics include: (1) precision (Pre), (2) recall (Rec), and (3) f-measure (F-score). Given the returned subset of nodes S and the corresponding true subset of anomalies S^*, we can calculate these metrics as follows:

$$\text{Pre} = \frac{|S \cap S^*|}{|S|}, \text{Rec} = \frac{|S \cap S^*|}{|S^*|}, \text{F-score} = \frac{2|S \cap S^*|}{|S^*| + |S|}$$

5.2 Semi-synthetic Data Evaluation

We evaluate the accuracy of our approach to detect the disrupted ground truth anomalies. Fig. 3 shows the average precision, recall, and F-measure in detecting the injected anomalies using the semi-synthetic data. We find that overall our approach consistently outperforms all other baseline methods.

Detection power. (1) **Our approach.** Our approach outperforms baseline methods especially at low perturbation intensities where the detection is harder to carry out, and the performance increases gradually with the increase of perturbation intensity. In particular, it has a high accuracy of detecting injected anomalies when $q_t \geq 10$ regardless of the size of injected anomalies. Measures of recall using $\text{NPHGS}_{\text{Twitter}}$ are as good as our approach while the other baseline methods are significantly worse than our approach especially when the size of disrupted ground truth anomalies is small and perturbation intensity is low. However, the measures of precision using $\text{NPHGS}_{\text{Twitter}}$ are much worse than our approach. Considering overall F-score, $\text{NPHGS}_{\text{News}}$ and $\text{NPHGS}_{\text{Twitter}}$ look similar to our approach when perturbation intensity is low while much worse than our approach when perturbation intensity is high. When we increase q_t, EventTree based methods perform worse than our approach, especially when the size of ground truth anomalies is small. (2) **NPHGS.** When $q_t \in \{1.0, 2.0\}$ and true ratio $x \in \{0.05, 0.10\}$, the precision of $\text{NPHGS}_{\text{News}}$ is better than our method. However, when $x = 0.15$, the recall of $\text{NPHGS}_{\text{News}}$ becomes quite low, which indicates its poor behavior when true subgraph is relatively large. (3) **EventTree.** The recall of $\text{EventTree}_{\text{News}}$ and $\text{EventTree}_{\text{Twitter}}$ is among the best when q_t is less than 2.0. The reason is that results of EventTree are easier affected by noise nodes.

Running Time. We compare the running time of our algorithm with baseline methods as shown in Fig. 4. The running time of our algorithm and NPHGS-based methods is proportional to the true anomalous subgraph size. However, the running time of ET-based methods are almost the same with respect to different true anomalous subgraph sizes as we tried a fixed number of λ on PCST algorithm. $\text{NPHGS}_{\text{LocalNews}}$ uses the least running time but gets the worst results. Overall, our algorithm is competitive to baseline methods in terms of running time.

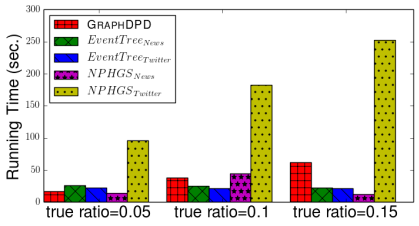

Figure 4: Comparison on running time.

5.3 Real Data Evaluation

We apply our proposed approach to TP and LNA of Mexico and Venezuela during Year 2014 as shown in Table 3. Performance evaluation on censorship detection is two-fold: (1) quantitative evaluation on accuracy, and (2) qualitative case studies.

5.3.1 *Quantitative Evaluation*

We apply our approach on graphs of all possible time windows from three days to seven days with starting days from January 1, 2014 to December 25, 2014, which are $5 \times 359 = 1795$ graphs in total. For each graph, we find the connected subgraph with the largest score as defined in Eqn. 2. We perform 5,000 random permutations and remove subgraphs whose p-values are greater than a predefined significance level (0.05). If a starting day has multiple significant subgraphs associated with different time windows, we just record the most significant subgraph with its corresponding time window for this day. We find 43 distinct significant subgraphs in Mexico and 54 distinct significant subgraphs in

Venezuela during Year 2014. In order to eliminate possible duplicated results, subgraphs are ranked based on p-values from low to high and removed if within three days of another subgraph with a lower p-value. After removing duplicates, we finally identify 12 distinct significant subgraphs in Mexico and 11 distinct significant subgraphs in Venezuela.

As discussed before, existing approaches on censorship detection in social media rely on the collection of deleted posts and such approaches are not capable to detect censorship in news media. Hence, we apply two anomaly detection baseline methods, NPHGS and EventTree, to detect anomalies in LNA on graphs of all possible time windows from three days to seven days with starting days from January 1, 2014 to December 25, 2014. The parameters used for the baselines are set similarly as in Section 5.1. The baseline methods can find the connected subgraph with the largest score in each graph according to their object functions, however, they are not able to evaluate the significance level of each subgraph. For the purpose of comparison, we rank the subgraphs detected by the baseline methods from the best to the worst according to their function values and compare top 12 subgraphs in Mexico and top 11 subgraphs in Venezuela with our method.

As there is no ground truth information of censorship, we formulate the following criteria and we believe a detected subgraph is an indicator of censorship if we find: 1) the incident referred by the cluster of keywords and associated time window is anti-government or a blame of government, 2) domestic newspapers in the target country did not report the incident around the time window associated with the subgraph, and 3) international newspapers reported the incident or verified the existence of censorship in domestic newspapers in the target country around the time window asscociated with the subgraph. Any detected subgraph fails to satisfy any of these criteria is considered as a false positive, which is detected as an indicator of censorship but we cannot find sufficient evidences. Table 4 summarizes the comparison of false positive rates in censorship detection and our method outperforms NPHGS and EventTree. The reason is that the baseline methods are designed for event detection instead of censorship detection. The cluster of keywords relevant to a censorship event should be burst in Twitter while silent in local news articles. The baseline methods can not differentiate censorship events from non-censored events, for instance, the end of attention on events. In contrast, our method is capable to distinguish true censorship patterns with the integration of Twitter data. Table 5 summarizes the top 4 censorship patterns we have detected in Mexico and Venezuela, which are ranked by their p-values from low to high. We will elaborate more details about how we verify the existence of censorship in the following section.

Country	NPHGS	EventTree	GraphDPD
Mexico	0.67	0.5	0.25
Venezuela	0.64	0.46	0.27

Table 4: Comparison of false positive rates in censorship detection between GraphDPD and two baseline methods: NPHGS and EventTree on real data of Mexico and Venezuela during Year 2014.

5.3.2 Case Studies

Mexico: May 2014. In December 2013, Mexican president Peña Nieto and Congress amended the Constitution, opening up the state controlled oil industry to foreign investors. Tens of thousands of protesters demonstrated in Mexico City on Labor Day (May 1) to protest against the energy reform, fearing the total privatization of the energy sector and higher energy prices [6]. In additions, protesters were also unsatisfied with the 2013 reforms of the educational sector. However, this incident was not reported in a number of influential newspapers in Mexico, including but not limited to Noroeste, Vanguardia, El Siglo de Torreon, Correo, El Imparcial, El Informador, Novedades Acapulco, and El Universal, which is an indicator of censorship. Fig. 5a shows a cluster of censored keywords detected by our method around May 1, 2014 in Mexico. Due to space limitations, Fig. 6 shows censorship patterns on a few censored keywords, however, other censored keywords have similar patterns as well. Our approach has successfully captured consistent censorship patterns among a collection of relevant keywords, which well describe the topics around which the May 1 demonstrations were organized (reforma, gasolina, dinero, educacion, escuela).

Venezuela: February 2014. As a result of the collapse of the price of oil (main export of the country), a decade of disastrous macroeconomic policies and growing authoritarianism Venezuela suffered from inflation, shortages of basic foodstuffs and other necessities, and increasing political frustration. Mass opposition protests led by opposition leaders demanding the release of the students occurred in 38 cities across Venezuela on February 12, 2014. The incident was reported by a number of major international newspapers such as BBC, CNN, and New York Times ([8, 10, 11]), and censorship in the country's largest daily Ultimas Noticias was confirmed in a number of international news outlets ([4, 5, 9]). The day after the protests President Maduro announced that Colombian TV news channel NTN24, which had been the only station to broadcast the protests to within Venezuela, was being removed from the grid of Venezuelan cable operators for airing anti-government demonstrations. Fig. 5b shows a cluster of censored keywords detected by our method around February 18, 2014 in Venezuela. Our approach has successfully captured consistent censorship patterns among a collection of relevant keywords, which well describes the populations involved (estudiante, chavistas, opositores, leopoldolopez), the target of the demonstrations (nicolasmaduro), and the reasons for the demonstrations (apoyo, heridos, libertad).

False positives. Our method identified a collection of keywords, which have consistent pattern of burst in Twitter posts while no observed significant changes in local news articles, during the time region from June 28th, 2014 to June 30th, 2014. Although the collection of keywords satisfy censorship pattern, they are actually relevant to soccer games in 2014 FIFA world cup instead of censorship given the detected keywords are only mentions of: game (fifa, fútbol, penalty, perder, pasar), soccer player or coach that are in the matches during the time period (miguel, herrera, james, rodríguez, robben), and countries that have matches during the time period (brasil, chile, costa, rica, holland, mexicano). Therefore, this is an example of false positive in our results.

6. CONCLUSION

In this paper, we have presented a novel unsupervised approach to identify censorship patterns in domestic news media using social media as a sensor. Through comprehensive experiments on semi-synthetic datasets, we showed that our approach outperforms popular anomalous subgraph detec-

Table 5: Top 4 censorship cases from our results in Mexico and Venezuela during Year 2014 (ranked by p-value from low to high)

Rank	p-value	Date	Example detected keywords	Why are relevant news censored?
Mexico				
1	0.00227	2014-05-01	reforma(reform), gasolina(petrol), educacion(education)	Tens of thousands of people marched in Mexico City on Labor Day to protest the new laws, which target at Mexico's education system and opening up the state controlled oil industry to foreign investors.
2	0.00379	2014-11-19	ayotzinapa, victimas(victims), normalistas, desaparecidos(missing), detenidos(detained), manifestantes(protesters), marcha(march)	Hundreds of thousands of people protested for the disappearance of 43 Ayotzinapa students at the Zocalo in Mexico City.
3	0.01258	2014-11-10	ayotzinapa, estudiantes(students), normalistas, desaparecidos(missing), protesta(protest), militares(military), iguala	Protests in Mexico City demanding the return of the missing students, who came from Ayotzinapa Rural Teachers' College and went missing in Iguala on September 26, 2014, turned violent for the first time. Protesters set fire at the National Palace and some of them were arrested.
4	0.01909	2014-11-27	murio(died), secuestro(kidnapping), estudiante(student), detenidos(arrested), falleció(died), qepd(RIP), fallecimiento(death), muerte(death)	Protesters were demanding justice for 43 students who were abducted and apparently murdered in September.
Venezuela				
1	0.002722	2014-12-24	mayor, cambio(change), lara, presos(prisoners), problema(problem), sigue(follow)	In the cities of Lara and Trujillo, students performed protests as those in Plaza Altamira calling for the release of student protesters who have been arrested since February.
2	0.00545	2014-02-18	represion(repression), disparo(shooting), marchamos(march), heridos(wounded), nicolasmaduro, armados(armed), leopoldolopez, apresar(arrest), ntn24	Mass protests led by opposition leaders, including Leopoldo López, occurred in 38 cities across Venezuela asking for the release of the arrested students. Colombian TV news channel NTN24 is taken off air for airing anti-government demonstrations.
3	0.00668	2014-05-01	muertes(deaths), cambio(change), caracas, presidente(president), labor	Thousands of Venezuelans demonstrated in Caracas to commemorate Labor Day and denounce shortages. Some protesters were injured when dispersed by authorities.
4	0.03243	2014-11-03	estudiantes(students), maracay, seguidores(followers), libertad(freedom), heridos(wounded), alcalde(mayor), psuv	Reference to events on February 18, 2014. Supporters of Leopoldo López were forcibly removed by police from where they were concentrated for seven days in Maracay.

(a) Mexico 2014-05-01 **(b)** Venezuela 2014-02-18

Figure 5: Word cloud representing censored keywords in LNA identified by our method

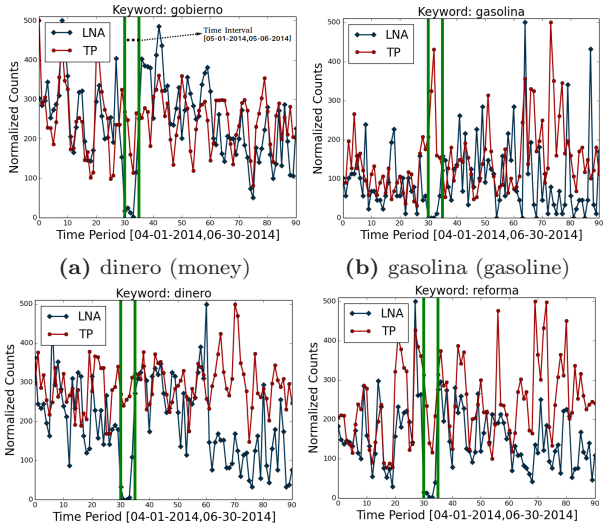

(a) dinero (money) **(b)** gasolina (gasoline)

(c) gobierno (government) **(d)** reforma (reform)

Figure 6: Example TSDF in LNA vs. TSDF in TP for a selected set of censored keywords detected by our method in Mexico. These keywords are relevant to the protests in Mexico City on Labor Day against education and energy reform. We can find consistent censorship pattern in LNA around Labor Day.

tion methods: EventTree and NPHGS, especially when the perturbation intensity is low. Analyzing real-world datasets in Mexico and Venezuela during Year 2014 demonstrates that our approach is capable of accurately detecting real-world censorship patterns in domestic newspapers. In future work, we are interested in generalizing censorship detection to other countries and to undertake censorship forecasting.

References

[1] Defending freedom of speech. http://saccityexpress.com/defending-freedom-of-speech/#sthash.cbI7lWbw.dpbs. Accessed on Jul 18th, 2016.

[2] Freedom House Worldwide Freedom of the Press 2014. https://freedomhouse.org/report/freedom-press/freedom-press-2014. Accessed on Jul 18th, 2016.

[3] Korea Policing the Net. Twist? It's South Korea. http://www.nytimes.com/2012/08/13/world/asia/critics-see-south-korea-internet-curbs-as-censorship.html. Accessed on Jul 18th, 2016.

[4] Once-defiant Venezuelan TV goes quiet amid opposition protests. http://articles.chicagotribune.com/2014-02-19/news/sns-rt-us-venezuela-protests-media-20140219_1_live-coverage-president-nicolas-maduro-news-channel-globovision/2. Accessed on Jul 18th, 2016.

[5] The Real Threat to Venezuela's Democracy. http://www.nybooks.com/daily/2014/04/09/venezuela-protests-censorship/. Accessed on Jul 18th, 2016.

[6] Thousands March In Mexico City On May Day. http://www.wbur.org/hereandnow/2014/05/02/may-day-mexico. Accessed on Jul 18th, 2016.

[7] Twitter Transparency Report. https://transparency.twitter.com/. Accessed on Jul 18th, 2016.

[8] Venezuela: At Least Two People Are Killed in Protests. http://www.nytimes.com/2014/02/13/world/middleeast/venezuela-two-people-are-killed-in-protests.html. Accessed on Jul 18th, 2016.

[9] Venezuela: Opposition Rallies End in Bloodshed, Riots. https://panampost.com/marcela-estrada/2014/02/13/venezuela-opposition-rallies-end-in-bloodshed-riots/. Accessed on Jul 18th, 2016.

[10] Venezuela student protest ends in deadly violence. http://www.bbc.com/news/world-latin-america-26166094. Accessed on Jul 18th, 2016.

[11] Venezuela: What's the crisis about? http://www.cnn.com/2014/02/20/world/americas/venezuela-qa/. Accessed on Jul 18th, 2016.

[12] Twitter Country Withheld Content Policy. https://support.twitter.com/articles/20169222, 2012. Accessed on Jul 18th, 2016.
[13] M. F. Alkazemi et al. Kuwaiti political cartoons during the arab spring: Agenda setting and self-censorship. *Journalism*, 16(5):630–653, 2015.
[14] A. A. Casilli et al. Social media censorship in times of political unrest-a social simulation experiment with the uk riots. *Bulletin of Sociological Methodology/Bulletin de Methodologie Sociologique*, 115(1):5–20, 2012.
[15] F. Chen et al. Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs. In *Proc. KDD*, pages 1166–1175. ACM, 2014.
[16] F. Chen et al. A generalized matching pursuit approach for graph-structured sparsity. In *Proc. IJCAI*, pages 1389–1395, 2016.
[17] A. Coskuntuncel. Privatization of governance, delegated censorship, and hegemony in the digital era: The case of turkey. *Journalism Studies*, pages 1–19, 2016.
[18] A. Di Florio et al. Bypassing censorship: a proven tool against the recent internet censorship in turkey. In *IS-SREW*, pages 389–394. IEEE, 2014.
[19] Y. Gao et al. Multimedia social event detection in microblog. In *MMM*, pages 269–281. Springer, 2015.
[20] C. Hegde et al. A nearly-linear time framework for graph-structured sparsity. In *Proc. ICML*, pages 928–937, 2015.
[21] A. Java et al. Why we twitter: understanding microblogging usage and communities. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, pages 56–65. ACM, 2007.
[22] T.-W. Kuo et al. Maximizing submodular set function with connectivity constraint: Theory and application to networks. *IEEE/ACM Transactions on Networking (TON)*, 23(2):533–546, 2015.
[23] D. B. Neill. An empirical comparison of spatial scan statistics for outbreak detection. *International journal of health geographics*, 8(1):1, 2009.
[24] S. Petrovic et al. Can twitter replace newswire for breaking news? 2013.
[25] N. Ramakrishnan et al. 'beating the news' with embers: forecasting civil unrest using open source indicators. In *Proc. KDD*, pages 1799–1808. ACM, 2014.
[26] A. Ritter et al. Open domain event extraction from twitter. In *Proc. KDD*, pages 1104–1112. ACM, 2012.
[27] P. Robinson et al. *Pockets of resistance: British news media, war and theory in the 2003 invasion of Iraq*. Oxford University Press, 2013.
[28] P. Rozenshtein et al. Event detection in activity networks. In *Proc. KDD*, pages 1176–1185. ACM, 2014.
[29] P. Seib. *Beyond the front lines: How the news media cover a world shaped by war*. Springer, 2016.
[30] I. Subašić et al. Peddling or creating? investigating the role of twitter in news reporting. In *ECIR*, pages 207–213. Springer, 2011.
[31] R. S. Tanash et al. Known unknowns: An analysis of twitter censorship in turkey. In *Proceedings of the 14th ACM Workshop on Privacy in the Electronic Society*, pages 11–20. ACM, 2015.
[32] H. Tong et al. Fast random walk with restart and its applications. In *Proc. ICDM*, pages 613–622. IEEE Computer Society, 2006.
[33] K. Watanabe et al. Jasmine: a real-time local-event detection system based on geolocation information propagated to microblogs. In *Proc. CIKM*, pages 2541–2544. ACM, 2011.

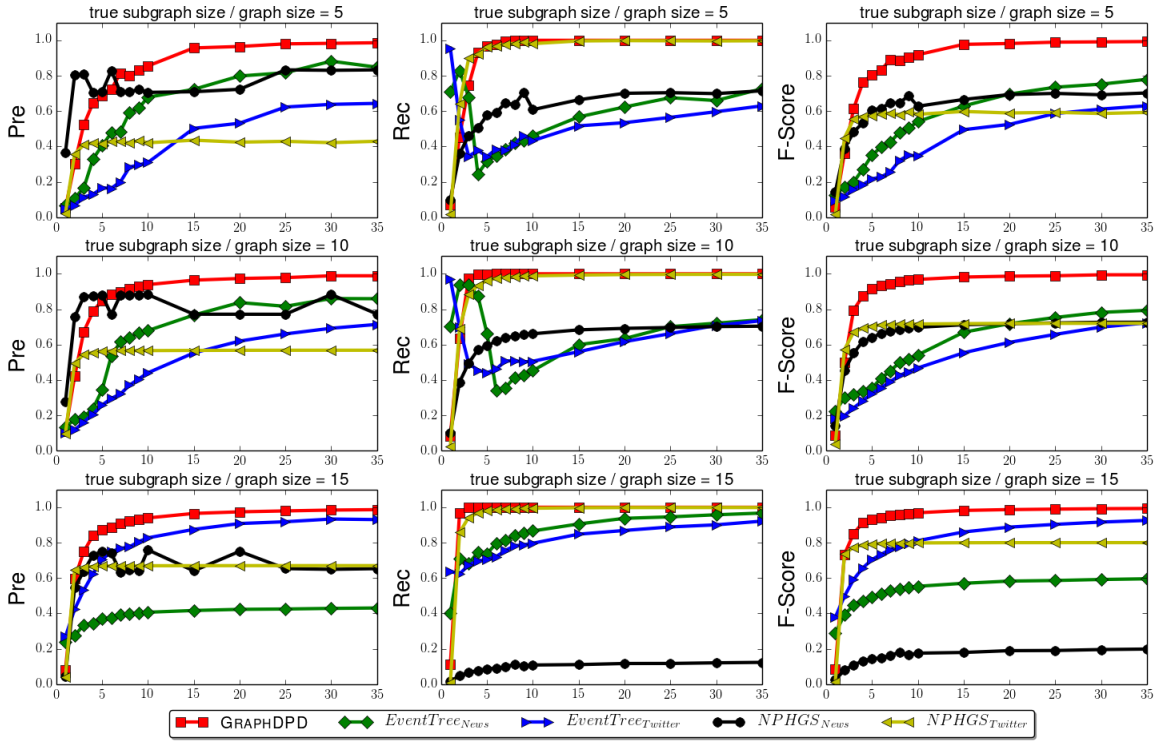

Figure 3: Anomaly detection results (mean precision (left), recall (center), and F-measure (right) vs. perturbation intensity) for different anomaly subgraph sizes (increased size from top to bottom) in semi-synthetic data. X-axis represents q_t, which implies the scale of anomaly injected in TP. q_n, which implies the scale of anomaly injected in LNA, is varied following $q_t \times q_n = 1.0$.